State Analytical Capacity and Consultocracy Through CPAF v6

Seven Lenses on How a State Loses, and Rebuilds, Its Capacity to Think — A Conceptual Proof of Concept for the Computational Policy Analysis Framework (CPAF)

Conceptual Working Paper for Discussion

Using CPAF's seven dimensions as an analytical typology for a new domain, with no model built or run

Bilel ABBOUDI

October 2026

Contents

Executive Summary 3

1. Strategic Framing and the Capacity Question 4

1.1 The Context 4
1.2 The Question Conventional Analysis Cannot Answer 5
1.3 The Methodological Gap 5

2. CPAF as a Typology: Seven Lenses on One Situation 5

2.1 The Lenses 5
2.2 How the Lenses Inform One Another 6
2.3 What the New Domain Required Changing 7

3. The Network Lens: Which Way Knowledge Flows 7

4. The Narrative Lens: Objectivity, Value for Money, and the Duty to Think 8

5. The Feedback Lens: The Loops That Hollow or Rebuild 8

5.1 The Key Quantities 8
5.2 The Central Loops 9
5.3 Delays, Absorptive Capacity, and the Disconfirming Evidence 10

6. The Actor Lens: Commissioners, Analysts, and the Revolving Door 11

7. The Configuration Lens: Where an Administration Could Settle 11

7.1 The Hollow-Dependent Configuration 12
7.2 The Capable-Owner Configuration 12
7.3 The Rehired Ghost 12
7.4 Shelfware: A Probably Transient Fourth Configuration 13
7.5 The Aggregate Masks the Structure; Asymmetry and Hysteresis 13

8. The Candidate Constraint and the Routes Out 13

8.1 Capacity and Method: The Central Hypothesis and Its Alternatives 13
8.2 Three Senses of Ownership 14
8.3 The Routes Out, by Leverage Depth 15
8.4 The Instrument Layer and the AI Caveat 16
8.5 Conditional Falsifiers 16

9. What This Shows, and the Next Stage 17

Annex A. Methodological Notes 17

A.1 Validation: What Transfers and What Does Not 17
A.2 What the Document Does and Does Not Claim 18
A.3 How This Document Was Produced 18

Annex B. References 19

A note on what this document is, and is not

This is a conceptual proof of concept and the analytical companion to the author's essay "Ghost in the Shell". Its purpose is to show that the seven dimensions of the Computational Policy Analysis Framework work as a typology of questions that anyone analysing a public-policy situation can bring to it, with or without building a model. Consultocracy, the state's reliance on consultants for its own thinking, is the test domain.

No model of this domain has been built or run. Everything here is a proposed mechanism, a hypothesised configuration or a candidate constraint; every position and magnitude is qualitative. Terms such as "attractor" and "loop" are used in their conceptual sense. Where a future computational build is discussed, it is identified as a separate, later stage.

The document represents mechanisms documented in the public-administration literature. It is not an assessment of any government, administration, firm or profession, a ranking of administrative traditions, a judgement on any reform programme, a prediction, or procurement advice.

***N.B.** References were checked against publisher or repository records where marked in Annex B.*

Executive Summary

Governments that have relied heavily on external consultants for their policy analysis are now trying to rebuild what they lost. Auditors recommend that the state invest in its own people and write knowledge transfer into contracts; parliamentary inquiries reach the same conclusion; several administrations have announced insourcing and in-house consulting units. Yet the consulting bill keeps rising, and the state often cannot say with confidence how much it spends. The question is no longer whether the state should think for itself, but why, once the habit of buying its thinking has formed, it is so hard to stop, and why some attempts to stop may simply change the badge on the same dependency.

This document brings the seven dimensions of the Computational Policy Analysis Framework (CPAF) to that question, not as software to be run but as a **typology**: seven lenses, each asking a different question of the same situation. Who is connected to whom, and which way does knowledge flow? Which stories make a choice look natural? Which loops reinforce the pattern and which balance it? Where could the system settle? How does the picture differ across places and over time? Whose decisions add up to it? And how confident are we in each answer? CPAF was developed as a computational framework in the author's master's research on Indian craft clusters (Abboudi, 2026). The claim here is narrower: the same seven questions organise a very different domain well, and an analyst, auditor or manager can use them without any modelling.

Applied to consultocracy, the lenses point to a **hypothesised set of configurations**. A state may settle into a **hollow-dependent** configuration, in which the offices are staffed but the framing, modelling and weighing of options are bought in. It may reach a **capable-owner** configuration, in which officials frame and choose and external experts augment their judgement. Or it may land in a third, less-remarked configuration this document calls the **rehired ghost**: capacity is brought back in-house, but the borrowed frames, methods and incentives come with it. The third is the specific danger the lenses make visible, because the two numbers most often used to track the problem, consultancy spend and in-house headcount, can both improve on the way into it. A fourth, **shelfware** (tools without the people to use them), is probably transient.

The **candidate constraint** is that leaving the hollow-dependent configuration takes both in-house analytical capacity and command of a method fit for public purpose. That is a hypothesis, not a

finding, and it has rivals: the two may reinforce each other gradually, or strength in one may partly compensate for weakness in the other. Each rival implies different advice, and only evidence from real administrations can separate them. The lenses also split the word "ownership" into three senses (control of assets and access, command of the method, fitness for public purpose), and the rehired ghost is the case where the first is restored and the second is not.

The evidence does not all point one way, and the document says so. The strongest quantitative study found in preparing it, a six-year panel of English public hospital trusts, found no sign that consultancy spending reduced later management headcount, and found that larger management functions used more consultants, not fewer; what predicted lower consultant use was the share of managers working in internal consulting functions (Kirkpatrick et al., 2023). This is consistent with the proposed mechanism, which concerns analytical function rather than headcount, but it is recorded as the most important disconfirming evidence and turned into a conditional falsifier in Section 8.5.

What the document shows is that CPAF's lenses can organise the questions and mechanisms of this domain. It does not show that any configuration is stable, that any constraint is necessary, or what will happen in any administration. Section 2.3 records what the new domain required changing in the framework, as well as what carried over.

1. Strategic Framing and the Capacity Question

1.1 The Context

The reliance is large and, in the best-documented case, still growing. In the United Kingdom, procurement data compiled by Tussell put the public sector's consulting bill at a record £3.7 billion in 2024/25, although that is still only about 2% of public procurement spending (Consultancy UK, 2026). In an industry survey, 96% of responding senior civil servants expected their departments' consulting spend to rise over the next two years (Consultancy UK, 2025). The National Audit Office found that the four data sources used in government to measure central-government consultancy spend differed by more than £1 billion for a single year, and concluded that there is no clear picture of what is spent or how it has changed (National Audit Office, 2025). The state cannot fully see its own dependency, which these lenses treat as part of the problem rather than a footnote.

The critical diagnosis is strong, though not uncontested. Heavy reliance on consultants monopolises and privatises public knowledge, erodes tacit knowledge inside the institution, weakens accountability, and strengthens a narrowly instrumental way of reasoning (Ylönen & Kuusela, 2019). Whether externalising policy advice produces a lasting loss of in-house capacity is, however, still an open empirical question: Veselý (2013) finds the evidence more varied than the universal-trend story suggests, differing across policy domains, countries and periods, and asks whether unused advisory capacity is lost or merely "sleeping". This document treats that question as its central hypothesis to test (Section 8.5). Comparative interview research shows the same firms playing opposite roles depending on whether in-house capacity survived: partners in formulation where the public service had been hollowed, bounded suppliers of specific expertise where it had not (Marciano, 2024).

The remedies on the table are also consistent: invest in the state's own people, use consultants only where they add genuine value, and write knowledge transfer into contracts (National Audit Office, 2025; Senate Finance and Public Administration References Committee, 2024). What is less often asked is how those remedies interact: which of them change the system and which only change the count.

1.2 The Question Conventional Analysis Cannot Answer

The instruments used to govern this problem operate at two levels. Spending data and spending caps track how much is bought; workforce data track how many people the state employs. Both are useful; neither answers the question that decides the outcome. Spend can fall while the thinking stays outside, moved to framework agreements, labour hire or vendor-held systems. Headcount can rise while the new posts administer contracts rather than do analysis, or while recruits from consulting carry its methods into permanent roles. The question that matters (is this administration positioned to keep its own judgement, or drifting towards dependency or a rehired ghost?) is a question about a system's trajectory, and neither number represents the system.

1.3 The Methodological Gap

There is a space between the spending ledger and the workforce return: a way of looking at the state's analytical capacity as one connected situation, taking in the network through which knowledge flows in or out, the narratives that make buying advice look objective, the loops of learning and forgetting, and the decisions of the officials who commission or do the work, and asking where that configuration is heading. CPAF's seven lenses are one way to occupy that space, whether or not a model is ever built. They do not replace audit, workforce planning or procurement control; they help ask the question none of them answers on its own.

2. CPAF as a Typology: Seven Lenses on One Situation

2.1 The Lenses

CPAF was built as a computational framework that couples seven analytical dimensions. Each dimension also stands on its own as a question worth asking of any public-policy situation, and the questions are more useful together than apart, because each catches something the others miss. Used this way, the framework is a typology: a structured checklist for analysis, not a program to run. The table sets out each lens, the general question it asks, and how it is asked of consultocracy in this document.

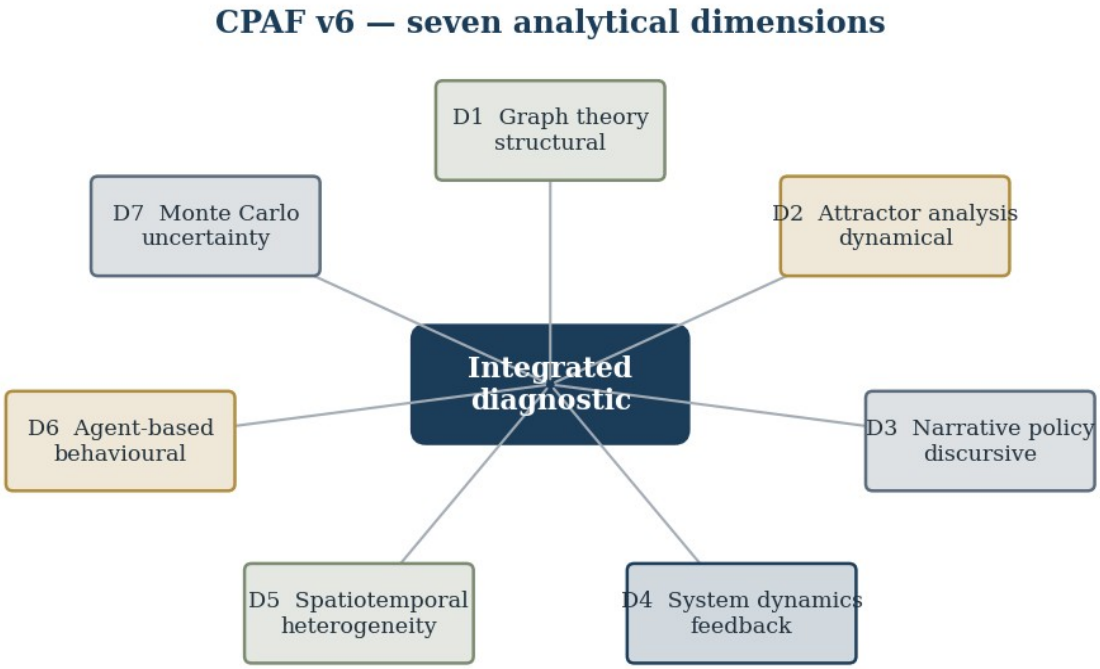


Figure. *CPAF's seven analytical dimensions, read here as seven lenses on one situation. The integrated diagnostic is what the lenses show together that none shows alone.*

Lens (CPAF dimension)	The question it asks of any situation	Asked of consultocracy
Network (D1, graph theory)	Who is connected to whom, how concentrated are the connections, and which way does what matters flow along them?	Which way does knowledge flow along each engagement: into the state, or out of it? (Section 3)
Configurations (D2, attractor analysis)	Where could the system settle, and what would it take to move it from one place to another?	Hollow-dependent, capable owner, rehired ghost, and a probably transient shelfware state (Section 7)
Narrative (D3, Narrative Policy Framework)	Which stories make one choice look natural and another look odd?	Objectivity, value for money, capability, accountability (Section 4)
Feedback (D4, system dynamics)	Which loops reinforce the current pattern, which balance it, and how long do changes take?	Learning displaced, memory drained, capability rebuilt, knowledge delivered inward (Section 5)
Place and time (D5, spatiotemporal analysis)	How does the picture differ across places, units and periods, and what lags separate cause from effect?	Differences across traditions and departments; capacity lost fast, rebuilt slowly (Sections 3, 5 and 7)
Actors (D6, agent-based modelling)	Whose individual decisions add up to the pattern, and how do they imitate one another?	Commissioning officials, analysts, recruits from consulting, firms (Section 6)
Uncertainty (D7, Monte Carlo)	How confident are we in each mechanism, and which conclusions survive if one is wrong?	Provenance declared per mechanism; rival hypotheses kept open (Section 8; Annex A)

2.2 How the Lenses Inform One Another

The lenses are not independent. A finding under one sharpens a question under another. Under the network lens, the direction of knowledge flow along an engagement tells the feedback lens whether that engagement builds capacity or draws it down. Under the narrative lens, the objectivity story tells the actor lens why a commissioning official reaches for a firm even when an in-house option exists. Under the actor lens, what recruits from consulting bring with them tells the feedback lens whether

insourcing adds capacity, method, or only headcount. Under the feedback lens, the speed of loss and rebuilding tells the configuration lens how costly it is to fall and how slow to climb back. And the uncertainty lens asks of every one of these links how well it is grounded. In the computational version of CPAF these links are implemented as couplings in code; here they are simply questions that should be asked in sequence.

2.3 What the New Domain Required Changing

A typology that fits every domain without change would be suspicious: it might be imposing its own shape. The table records what carried over from the origin application, what had to change, and where the fit is weak. Readers should weigh the recurring patterns (a combination constraint, several possible configurations) with this in mind: they may partly reflect the template brought to the analysis.

In the origin application (craft clusters)	In this domain	What had to change
The outcome was a measurable household quantity (income), with a natural unit	The outcome is the state's analytical capacity, which has no natural unit	Capacity is defined by function (people doing analytical work) and flagged as a proxy needing evidence of competence, time and independent judgement
The trap ran through market structure: a dominant intermediary between producer and buyer	No single supplier controls the state; concentration exists but is not the core mechanism	The trap mechanism changed: from market power to learning displaced inside the state. The network lens had to carry edge direction, not only concentration
Knowledge and skill decayed in households and clusters	Knowledge can flow in and be displaced at the same time, along the same engagement	Two separate flows were needed (knowledge delivered inward; learning displaced), where the origin needed one
No counterpart to a professional ethos	Public-service ethos decides what capacity is used for	A new element, treated as a moderator of what capacity produces, not as a stock that adds to capacity
Actors were artisans, intermediaries and agencies	Actors include recruits who move between firms and the state	A new mechanism, carry-over of method and ties through the revolving door, with no origin analogue
Place meant geographic clusters with distinct conditions	Place means administrative traditions, departments and levels of government	The spatial lens fits weakly here; its temporal side (lags, asymmetry) does most of the work
Uncertainty was propagated through a running model	There is no running model	The uncertainty lens becomes a discipline of declaring provenance and keeping rival hypotheses open
A combination constraint appeared in the origin analysis	A combination of capacity and method is proposed here	Kept as one of three rival hypotheses, because a framework that has seen a gate once may look for one again

3. The Network Lens: Which Way Knowledge Flows

Consultocracy is a network phenomenon before it is a budget line. The nodes are departments, the firms that advise them, and the individuals who move between the two; the edges are contracts, standing “strategic alliance” arrangements, and the revolving door through which senior officials move into private-sector roles (Marciano, 2024; Andrews & Beynon, 2024). The network lens asks what a contract-by-contract review cannot: how concentrated the state's advice is in a few suppliers, which departments sit at the centre of that dependency, and, most importantly, the direction in which knowledge travels along each edge.

That last question gives the essay's central distinction a definite shape. An engagement in which a subject-matter expert supplies bounded knowledge and leaves it behind is an edge along which

knowledge flows **into** the state; an engagement in which a firm performs the analytical function and returns a finished product is an edge along which the capacity to do that work flows **out**. The same firm, and even the same contract value, can sit on either kind of edge. The difference is structural role, not competence, which is why the lens asks about edge direction rather than supplier identity.

Two diagnostics follow. **Concentration analysis** asks whether the advisory network has a few highly central suppliers on whom many departments depend, the network signature of a dependency that is hard to unwind. Such concentration is documented: an audit cited by Howlett and Migone found that twenty-five suppliers accounted for 37% of UK departments′ spending on professional services (Howlett & Migone, 2013). The revolving door adds a second seam, along which conflicts of interest form (Asai, Kawai & Nakabayashi, 2021). **Flow-direction analysis** asks about the balance of inward and outward edges for each department, a reading of whether its external engagements are building its capacity or drawing it down. Both feed the feedback lens, because edge direction decides whether an engagement adds to the capacity stock or subtracts from it.

4. The Narrative Lens: Objectivity, Value for Money, and the Duty to Think

Whether a department buys its analysis or does it depends not only on budgets but on the account of external advice that circulates among those who commission it. Following the Narrative Policy Framework (Jones et al., 2023), this lens treats these accounts as policy-relevant in their own right. At least four contend. An **objectivity** narrative presents the external firm as neutral and the internal analyst as compromised. A **value-for-money** narrative presents buying as cheaper and faster than building. A **capability** narrative holds that the public sector needs dynamic capabilities of its own (Kattel & Mazzucato, 2018), and that a loss of knowledge undermines the state′s capacity to govern its relationships with the private sector (Mazzucato & Collington, 2023). And an **accountability** narrative, strengthened by public scandal, holds that judgement exercised outside government escapes the duties that bind public officials (Senate Finance and Public Administration References Committee, 2024).

The pivot of the domain is the objectivity claim, because it is the claim that the incentive structure quietly reverses. Where the analyst′s next contract depends on the client′s satisfaction, the temptation to produce policy-based evidence (evidence sought to support plans already made: Parkhurst, 2017; a transformation whose mechanisms are socially and institutionally embedded: Strassheim & Kettunen, 2014) stops being an individual failing and becomes structural. That extension to the commercial incentives of external analysts is this document′s own reasoning, not a finding of either source. The question this lens asks is one of coherence: does the objectivity narrative still hold within an administration, and is the capability narrative strong enough to make doing the work in-house the default rather than the exception? The answer conditions the feedback lens: the same budget goes to a firm or to an in-house team depending on which story the commissioning official believes.

5. The Feedback Lens: The Loops That Hollow or Rebuild

5.1 The Key Quantities

Five quantities carry the domain. Their definitions draw on the policy-capacity literature, which distinguishes analytical, operational and political capacity at individual, organisational and systemic levels (Wu, Ramesh & Howlett, 2015); this document concentrates on analytical capacity at the organisational level.

- **In-house analytical capacity (K)**: the state′s ability to frame its problems, model its options and weigh them on its own analysis. Read by function, not headcount: the people actually doing that work. Even that is a proxy. A fuller measure would need evidence of competence, of the time actually available for analysis, and of independence of judgement.
- **Outsourced share of analytical work (S)**: the proportion of framing, modelling and option appraisal performed outside the state and returned as a product.
- **Tacit institutional memory (M)**: the uncodified knowledge of what was tried, why it failed and who knows what (Polanyi, 1966/2009). Depleted by turnover; slow to rebuild.
- **Documented, usable instruments (I)**: models, data, methods and documented reasoning that the state can access, understand and audit, whether it owns them or holds the rights it needs. Instruments hold part of what is learned. They do not capture all tacit knowledge, so they reduce the loss caused by turnover; they do not remove it.
- **Public-service ethos (E)**: the commitment to accountability under law, equal treatment and a public interest no market prices (Moore, 1995; Bozeman, 2007), learned inside public institutions over time. E does not add to capacity. It moderates what capacity is used for, and in particular whether the instruments it produces fit public purpose.

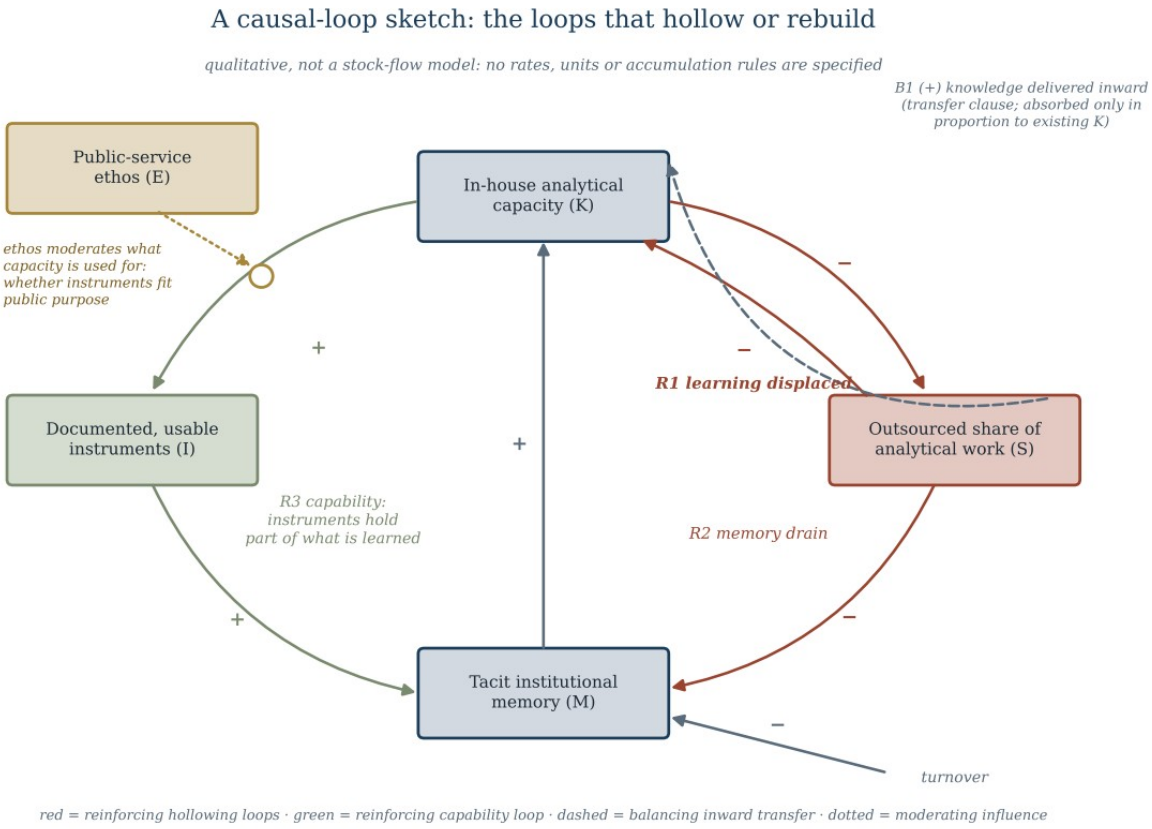


Figure. *A causal-loop sketch, not a stock-flow model. R1 (learning displaced) and R2 (memory drain) are reinforcing loops that deepen dependency; R3 (capability) is the reinforcing loop that rebuilds it; B1 is the balancing loop of knowledge delivered inward, absorbed only in proportion to existing capacity. Ethos moderates the link from capacity to instruments. A computational version would additionally need accumulation rules, rates and units for each quantity.*

5.2 The Central Loops

R1, learning displaced. Low in-house capacity raises the outsourced share; a higher outsourced share removes the learning-by-doing through which capacity is built; capacity falls further. This is the mechanism the essay calls “the dependency that feeds itself”, named a “vicious cycle” by a senior

public servant interviewed by Marciano (2024), described in the critical literature as short-term, outsourced knowledge production replacing the long-term work of civil servants (Ylönen & Kuusela, 2019), and identified by Mazzucato and Kattel as outsourcing that displaces learning-by-doing and hollows out cumulative expertise (Mazzucato & Kattel, 2026).

B1, knowledge delivered inward. External engagements can also transfer knowledge back, and contract clauses can require it (National Audit Office, 2025). This is a separate flow from R1, and the two run along the same engagement at the same time: a supplier can deliver knowledge inward while the opportunity to learn by doing is displaced. An engagement can therefore score well on transfer and still weaken capacity, if staff receive knowledge they never practise using. B1 is balancing but probably weak, because the supplier's commercial incentive does not favour making itself unnecessary. Its strength is unknown; it is treated as a range, and would be a priority for any planned sensitivity analysis.

R2, memory drain. Outsourcing and turnover deplete tacit memory; with less memory, each new task starts closer to zero and is more likely to be bought in. Talent routing may reinforce it: consulting has long been a sought-after destination for many of the ablest graduates (Mazzucato & Collington, 2023), and the literature records concern that lucrative consulting careers can dissuade those drawn to a public vocation (Seabrooke & Sending, 2022). How far this diverts talent away from public service is not measured here and is treated as an assumption.

R3, capability. In-house work builds capacity, and capacity used well produces documented instruments that hold part of what is learned, so that some of what one analyst knows is still available after they leave. R3 is the loop any remedy must strengthen. Because instruments hold only part of what is learned, R3 reduces the damage turnover does; it does not make capacity immune to it.

5.3 Delays, Absorptive Capacity, and the Disconfirming Evidence

Delays are differentiated. Capacity can be cut within a budget cycle; tacit memory and analytical capability accumulate over years. Different rates of loss and rebuilding establish **asymmetry**, which already makes sequencing matter: cutting external support before in-house capacity exists creates a gap that is filled, under pressure, by buying again. Asymmetry does not on its own establish **hysteresis**, the stronger claim that restoring earlier conditions would not restore the earlier state. That remains a hypothesis (Section 7.5).

A prerequisite links the loops. An organisation's ability to recognise, assimilate and apply external knowledge depends on the related knowledge it already holds, its absorptive capacity (Cohen & Levinthal, 1990), a concept applied directly to public services by Harvey et al. (2010). Under this lens, inward transfer is conditioned on existing capacity: a hollowed department absorbs less of what a good adviser leaves behind. Absorptive-capacity research motivates that dependence; it does not settle its form. Making transfer proportional to capacity is one choice among several (a threshold, or diminishing returns, would also fit the literature), and a future build would need to test the alternatives.

The strongest quantitative evidence cuts against the simplest version of R1. In English public hospital trusts over six years, larger management functions were associated with more consultancy spending, not less, and consultancy spending did not predict later reductions in management numbers (Kirkpatrick et al., 2023). The same study found that a higher share of managers working in internal consulting functions was associated with lower consultancy spending. This document does not explain the finding away. It records three things: R1 concerns analytical function, not headcount, so a

no-substitution result at headcount level does not test it; the association between internal consulting roles and lower external use is consistent with a capacity-driven mechanism; and Section 8.5 states the conditions under which evidence on analytical function would count against R1.

6. The Actor Lens: Commissioners, Analysts, and the Revolving Door

An administration's dependency is the aggregate of many decisions by individual officials: to commission work or do it, to stay or to move, to bring a method in or leave it at the door. The actor lens represents that variety. The actors include career analysts, commissioning managers, recruits from consulting, and the firms themselves. The empirical record warns against treating officials as passive: public-sector managers actively, and sometimes politically, drive consultancy use (Kirkpatrick et al., 2023), and junior officials commission work from firms they already know (Marciano, 2024).

Two questions follow. **Norm-tipping**: is "reach for the firm" or "do it in-house" the self-reinforcing norm among commissioning officials, through imitation? Once one norm dominates, new budgets and new tasks bounce off the other. **Carry-over**: what does and does not travel when a person moves from a firm into the state? The essay names three things: the imported framework and the residual tie to the former firm, which need examining, and the public-service ethos, which needs building. The literature documents a wider shift in which managers increasingly work as consultants do (Sturdy, Wright & Wylie, 2015). Under this lens, each recruit adds to capacity but may add to ethos only slowly, and may bring the borrowed method into the instruments. Where the method is examined and the ethos is built, insourcing rebuilds capacity; where they are not, insourcing may lead towards the rehired ghost. Mimetic adoption (organisations modelling themselves under uncertainty on others they perceive as more legitimate or successful, with a small set of major consulting firms spreading a few organisational models "throughout the land": DiMaggio & Powell, 1983) is a candidate mechanism by which that configuration could persist.

7. The Configuration Lens: Where an Administration Could Settle

The configuration lens turns the others into a strategic question: not how each contract performs, but where an administration's analytical capacity is heading under its current arrangements. Together the lenses suggest at least three configurations an administration could settle into, and a fourth that is probably transient. All four are hypothesised. No domain model has been run to establish that any of them is stable, and "attractor" is used here in its conceptual sense.

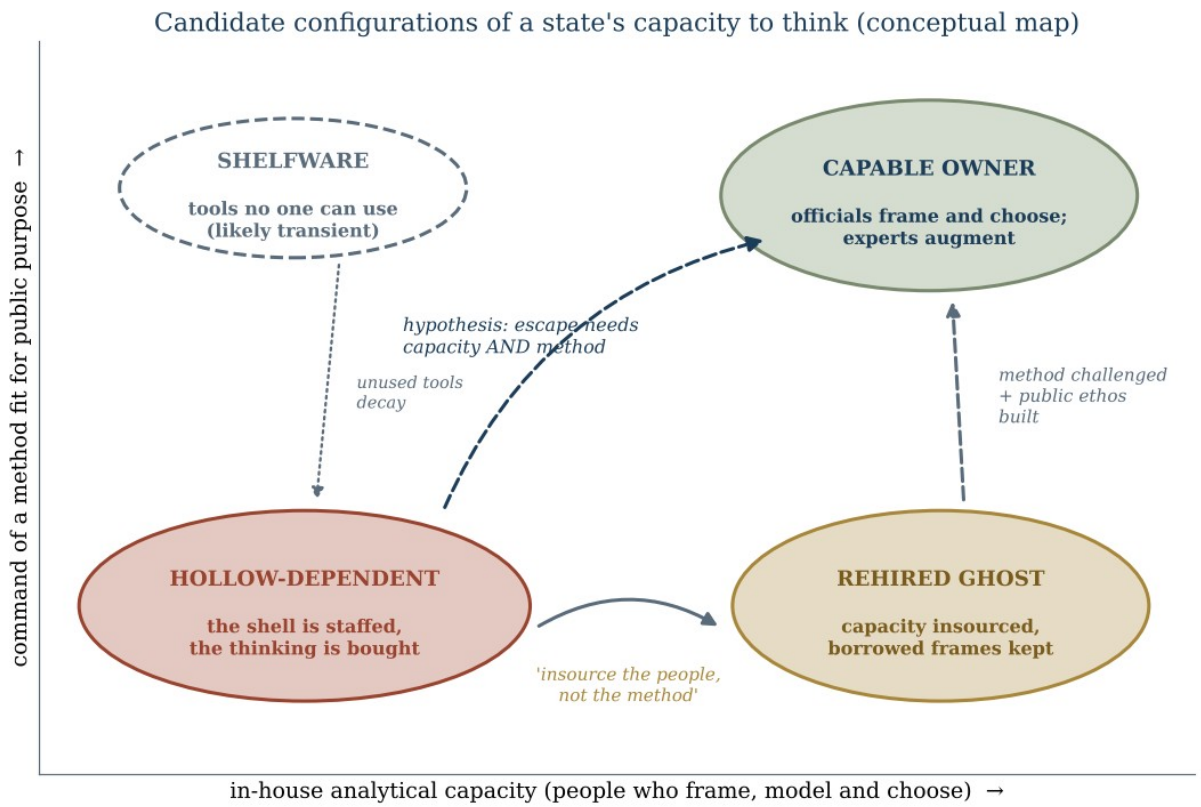


Figure. *Candidate configurations (conceptual map, not a measured model). Hollow-dependent (low capacity, weak command of method) is the hypothesised trap; capable owner (both strong) the escape; the rehired ghost, capacity brought back in-house while the borrowed method is kept, a settled but misleading configuration. Shelfware, tools without the people to use them, is shown dashed because it is probably transient.*

7.1 The Hollow-Dependent Configuration

Analytical capacity is thin, tacit memory is depleted, and the core work of framing and choosing is performed outside the state. Each engagement is individually rational, because the department genuinely cannot do the work this quarter, and together they keep the department unable to do it next quarter. If the proposed loops dominate, the configuration maintains itself: R1 and R2 both run in the hollowing direction, and B1 is too weak to reverse them.

7.2 The Capable-Owner Configuration

Officials frame problems and choose among options on their own analysis; documented instruments hold part of the reasoning; external experts are engaged for bounded knowledge and leave it behind. Consultants have a real and legitimate role here, the role the comparative evidence documents where in-house capacity survived (Marciano, 2024). R3 dominates. Turnover still costs something, but less, because part of what is learned is held in instruments as well as in people.

7.3 The Rehired Ghost

The third configuration deserves emphasis because current metrics cannot see it, and because it is where a well-intended remedy can quietly land. Capacity has been brought back in-house (posts are filled, an internal unit exists, spending may have fallen) but the method, the frames and the incentives of the consulting model have come with it. The internal team produces the same decks, maturity models and monetised business cases, applied to public problems whether or not public value survives their assumptions. It is not a failure on any spending or headcount measure, which is exactly why it is dangerous: it is the hollowing made invisible because it now looks solved. In the terms of

Section 8.2, the first sense of ownership (control of assets) has been restored and the second (command of the method) has not. Whether and when this configuration is acceptable is a choice an administration should make explicitly. This document makes no assessment of any specific insourcing programme, all of which are recent and whose trajectories are not yet observable.

7.4 Shelfware: A Probably Transient Fourth Configuration

Instruments without the capacity to use and audit them (a model no one can run, a dashboard no one questions) form a fourth configuration. It is shown dashed on the map and kept off the main argument because it is probably not stable: unused instruments decay, and the administration drifts back towards the hollow-dependent configuration, or into a new dependency on whoever maintains the tools. Whether it can persist, for example under a long maintenance contract, is an open question.

7.5 The Aggregate Masks the Structure; Asymmetry and Hysteresis

Two structural properties close the diagnostic. First, the configurations are not distinguishable by the headline numbers: spend can fall and headcount can rise identically on the way to the capable owner and on the way to the rehired ghost. The lenses therefore point to other readings: the direction of knowledge flow, the outsourced share of analytical work, and command of the method. Second, the speed of change differs by direction. What the evidence supports is asymmetry: a department can be hollowed within a few budget cycles, while rebuilding tacit memory and analytical capability takes years. That alone makes prevention cheaper than cure and the order of remedies as important as their content. The stronger claim, that the system has more than one domain of attraction (Holling, 1973) and that transitions between them show hysteresis, so that restoring earlier conditions is not enough to reverse the shift (Scheffer, 2009), is a hypothesis to test, not a finding.

8. The Candidate Constraint and the Routes Out

8.1 Capacity and Method: The Central Hypothesis and Its Alternatives

The central hypothesis of this document is that restoring a state′s capacity to think requires two conditions together: **in-house analytical capacity**, meaning people who frame problems, model options and weigh them themselves; and **command of a method fit for public purpose**, meaning a method those people can examine, choose between and adapt. The public-service ethos is what ties both to the public interest.

Each condition without the other fails in a recognisable way. Capacity without command of method produces the rehired ghost: analysts on the public payroll applying a borrowed method they have not examined. Instruments without capacity produce shelfware, or a new dependency on whoever maintains them. Recent theoretical work on public-sector capabilities reaches a parallel conclusion: public value emerges from the interaction of structural capacities, organisational routines and dynamic capabilities, not from any single attribute (Mazzucato & Kattel, 2026).

What remains open is how strictly the two conditions depend on each other. The strongest version holds that both are necessary, and that version cannot be established by modelling alone: a model built on the assumption that escape requires both conditions will, by construction, show that only their combination succeeds, and so confirm what was put into it. This document therefore states the hypothesis in three versions. They differ in one respect only: **how far strength in one condition can make up for weakness in the other.**

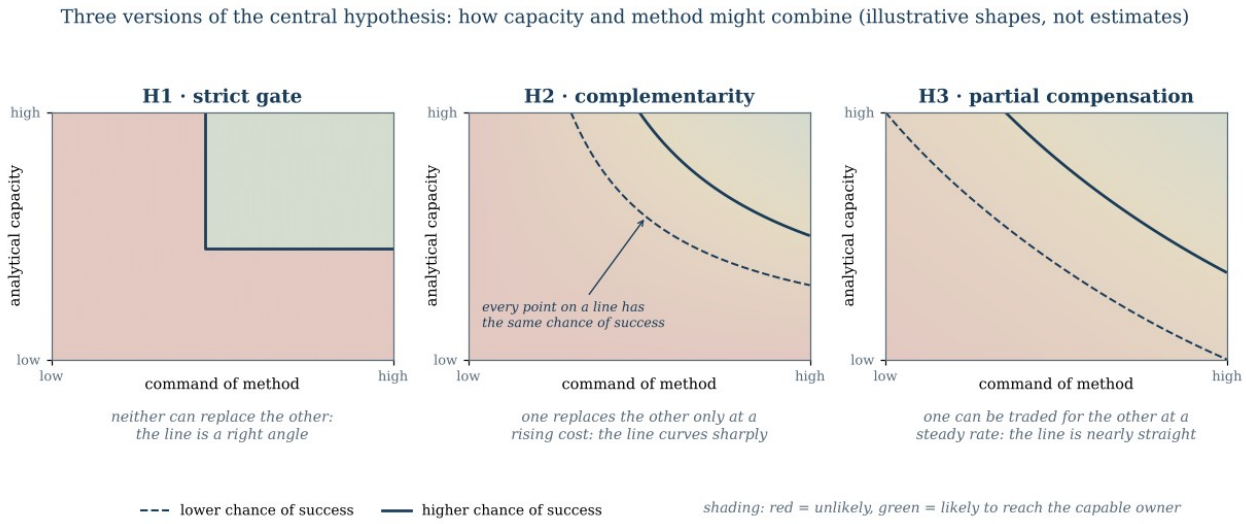


Figure. *Three versions of the central hypothesis. Each panel places command of method on the horizontal axis and analytical capacity on the vertical axis, from low to high. Shading shows how likely an administration is to reach the capable-owner configuration, from unlikely (red) to likely (green). Each line joins combinations that are equally likely to succeed (the dashed line a lower chance, the solid line a higher one); in H1 the two coincide at the threshold. The shape of the line carries the argument: a right angle means neither condition can replace the other, a sharp curve means one can replace the other only at a rising cost, and a near-straight line means one can be traded for the other at a steady rate. The shapes are illustrative, not estimates.*

- **H1, strict gate (no substitution).** Each condition is necessary, and below a minimum level of either, no amount of the other helps. The line forms a right angle: moving along one axis without the other changes nothing. *Implication: partial investment in one condition alone is largely wasted; the two must be built together.*
- **H2, complementarity.** The two conditions are complements: each raises the value of the other. An additional analyst is worth more where the method is well understood, and a well-understood method is worth more where there are analysts to use it. Neither is strictly necessary, but either alone achieves little. The lines curve sharply towards the axes: an administration weak in one condition can compensate only by becoming much stronger in the other, and beyond a point not at all. *Implication: progress on one condition is worth having, but its return depends on the other following.*
- **H3, partial compensation.** The two conditions are partial substitutes: strength in one can stand in for weakness in the other, at a roughly steady rate and for a time. A unit with few analysts can rely on secondments, a university partner or well-documented tools while it recruits; a unit with strong analysts can work with imperfect tools while it improves them. The lines are close to straight: losing some of one condition can be offset by gaining a comparable amount of the other. *Implication: starting with whichever condition is cheaper or quicker to build is a sensible strategy.*

The three versions lead to different advice on where to put the next budget, which is why the distinction matters in practice. They cannot be separated at a desk. Doing so would require evidence from administrations, for example comparing units that rebuilt their analytical teams with and without a deliberate change of method, on measures of the quality and independence of their analysis rather than on spending or headcount.

8.2 Three Senses of Ownership

The lenses also show that "ownership" covers three different things, and that an administration can succeed on one and fail on another.

- **Control of assets and access**: the ability to inspect, use, reuse, adapt and continue data, models and files. This does not always require owning them; UK government guidance notes that the right to use intellectual property is often more important than ownership, and that rights should be considered case by case (Digital, Data and Technology Profession, 2023).
- **Command of the method**: the ability of staff to examine a method, choose between it and alternatives, and adapt it when circumstances change.
- **Fitness for public purpose**: whether the analysis answers a public question, with the people affected and the outcomes that matter to them in view.

A government may own a model and apply it uncritically; it may use a licensed method and keep real independence of judgement. The rehired ghost is the case where the first sense is restored and the second is not. Throughout this document, the "method" in the candidate constraint means the second and third senses, not only the first.

8.3 The Routes Out, by Leverage Depth

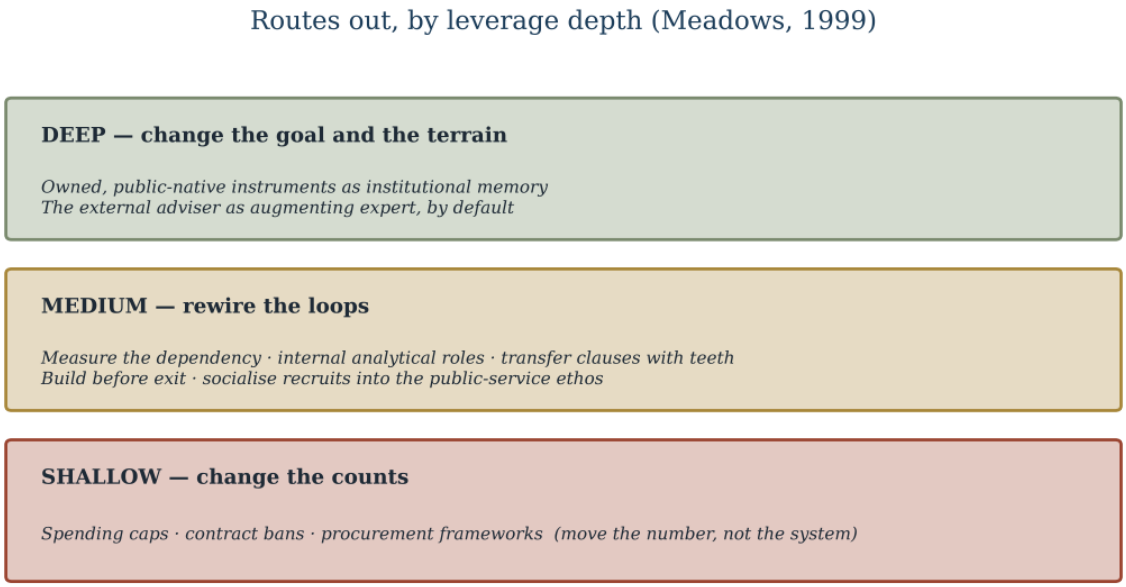


Figure. *Routes out of the hollow-dependent configuration, ordered by leverage depth (Meadows, 1999). Shallow moves change the counts; medium moves rewire the loops; deep moves change the goal and the terrain.*

Shallow moves change the counts. Spending caps, contract bans and procurement frameworks move the headline number without changing the configuration. On their own they risk displacing the dependency into other contract forms, or opening a service gap that is filled by buying again. They are useful as signals and as brakes, not as the remedy.

Medium moves rewire the loops. Five follow from the proposed mechanisms.

- **Measure the dependency first.** Track the outsourced share of analytical work and the direction of knowledge flow, not only spend. A state that cannot see its dependency cannot govern it (National Audit Office, 2025).
- **Build internal analytical roles, not just headcount.** The quantitative evidence links lower external use to managers in internal consulting functions, not to the size of the management function (Kirkpatrick et al., 2023).

- **Give knowledge-transfer clauses teeth.** Make transfer a deliverable with acceptance tests agreed in advance and run before final handover, and pair it with the in-house capacity to absorb it, since transfer into a hollowed department does little (Cohen & Levinthal, 1990).
- **Build before exit.** Sequence the rebuilding of capacity ahead of the withdrawal of external support, because of the asymmetry in Section 7.5.
- **Socialise recruits into the public-service ethos.** Treat the building of ethos as a process with its own pace, and examine the method recruits bring, so that insourcing people does not import the borrowed method by default.

Deep moves change the goal and the terrain. The deepest move is to make documented, inspectable, public-native instruments the state's institutional memory: models and methods built on named theory, documented, auditable, and either owned or held with full rights to use, adapt and continue them, so that what is learned survives turnover and contract cycles. The paradigm shift beneath it is to make the external adviser an augmenting expert by default, engaged for bounded knowledge the state cannot economically hold and required to leave it behind, so that the capable-owner configuration is the design rather than the exception.

8.4 The Instrument Layer and the AI Caveat

Artificial intelligence bears on this domain in two opposite ways, and the lenses above show why. Under H3, it could act as a compensating instrument: a unit with few analysts, using AI tools it can inspect and verify, may carry out analytical work it could not otherwise do while it rebuilds its team. Under R1, it could become a new channel of displacement: if the tool frames the problem and drafts the answer, officials lose the practice through which capacity is built, exactly as they do when a firm does the work. Which effect dominates depends on the three senses of ownership in Section 8.2: whether the state controls access to the tool and its outputs, whether its staff can examine and challenge how the tool reached its result, and whether the questions put to it are public ones. A closed system bought as a service fails the second test however well it performs, and occupies the same structural role as the consultant, with a different supplier.

This document therefore treats AI not as a remedy in itself but as an accelerant of whichever loop already dominates. The evidence on AI in public analytical work is recent and changing quickly; this section states a hypothesis to test, not a finding. A fuller treatment is the subject of the author's separate proof of concept on AI in public administration.

8.5 Conditional Falsifiers

Four kinds of observation would count against the proposed mechanisms, under the conditions stated. They are recorded before any domain build.

- **Learning displaced (R1).** If measures of analytical function, not headcount, show no erosion as the outsourced share rises, in units where recruitment, training and other routes to learning did not rise to offset it, R1 would be contradicted. Stable capacity alongside rising outsourcing is not enough on its own: offsetting learning has to be ruled out first.
- **The candidate constraint.** If units that rebuild capacity without examining their methods reach the same outcomes, on measures of independent judgement rather than spend and headcount and over a period long enough for differences to show, as units that also build command of method, the strict gate and the rehired-ghost configuration would be contradicted. If partial investment in either one shows proportional gains, the evidence favours H2 or H3 over H1.

- **Absorptive capacity.** If knowledge-transfer clauses raise capacity equally in hollowed and capable units, comparing engagements of similar scope and similar quality of transfer, the dependence of transfer on existing capacity would be contradicted.
- **Asymmetry and hysteresis.** If rebuilding proves as fast as losing, the asymmetry would be contradicted. If restoring the earlier conditions (budgets, posts, roles) restores the earlier capacity within a comparable time, hysteresis would be contradicted even if the asymmetry holds.

9. What This Shows, and the Next Stage

This proof of concept shows one thing: that CPAF's seven lenses can organise the questions and mechanisms of a domain very different from the one they were developed on. Applied to the state's own analytical capacity, they surface a distinction the usual numbers miss (the rehired ghost), a candidate constraint with explicit rivals, and a sharper meaning of ownership. Section 2.3 records what had to change. The recurrence of a combination constraint and of several configurations may partly reflect the template brought to the analysis, and is offered for testing, not as confirmation.

For practitioners, no model is needed. The lenses can be used as a checklist of questions on any situation (Section 2.1), and the companion practitioner templates, *Before You Sign*, translate them into decisions a manager can take on a real consulting proposal.

A computational build would be a separate, later stage, and a defined piece of work rather than a matter of tuning. It would need a stock-flow specification with accumulation rules, rates and units for each quantity and an explicit treatment of the outsourced share; a map of an administration's advisory network with the direction of knowledge flow on each edge; an actor layer parameterised from interviews with commissioning officials; alternative specifications for the candidate constraint and for absorptive capacity; and a validation suite run afresh for this domain. The natural interlocutors are the public-administration, audit and civil-service-reform communities, for whom the value is not a score but a structured account of where a configuration may be heading and what could move it.

Annex A. Methodological Notes

This annex records the status of the work. It is not required for engagement with the argument.

A.1 Validation: What Transfers and What Does Not

The parent framework has a validation suite organised on Barlas's three stages: direct structure tests, structure-oriented behaviour tests and behaviour-pattern tests, each contributing a different kind of evidence (Barlas, 1996). That suite is being re-run on the framework's rebuilt engine, and none of it has been applied to this domain, because no domain model exists to test. Test procedures could be transferred to a future build; passing status could not, because the variables, their meanings and the causal relationships have changed (Section 2.3). No validation claim is made here. Provenance is therefore declared per mechanism.

Mechanism	Principal source	Status
Same firms, opposite roles, depending on surviving in-house capacity	Marciano (2024)	Empirical: two-case comparative interviews
Internal consulting roles associated with lower external use; no headcount substitution	Kirkpatrick et al. (2023)	Empirical: one country, one sector, 2008–14

Mechanism	Principal source	Status
State cannot measure its own dependency	National Audit Office (2025)	Empirical: one central government
R1 learning displaced	Ylönen & Kuusela (2019); Mazzucato & Kattel (2026)	Literature-grounded; contested at headcount level (Kirkpatrick et al., 2023) and in the externalisation evidence (Veselý, 2013)
B1 knowledge delivered inward, separate from R1	National Audit Office (2025)	Mechanism-grounded; strength unknown, treated as a range
R2 memory drain and talent routing	Polanyi (1966/2009); Mazzucato & Collington (2023); Seabrooke & Sending (2022)	Mechanism-grounded; talent diversion assumed
Transfer conditioned on existing capacity	Cohen & Levinthal (1990); Harvey et al. (2010)	Mechanism-grounded; functional form open
Ethos as moderator of what capacity is used for	Moore (1995); Bozeman (2007)	Mechanism-grounded; moderating form assumed
Incentive towards policy-based evidence	Strassheim & Kettunen (2014); Parkhurst (2017)	Mechanism-grounded; extension to consultants is this document's reasoning
Rehired ghost through mimetic adoption and consultant-managers	DiMaggio & Powell (1983); Sturdy, Wright & Wylie (2015)	Mechanism-grounded; persistence hypothesised
Capacity–method combination (H1, H2, H3)	Mazzucato & Kattel (2026) as a parallel	Hypothesis with explicit rivals
Rights versus ownership	Digital, Data and Technology Profession (2023)	Guidance; the three-sense distinction is this document's
Asymmetry of loss and rebuilding	Polanyi (1966/2009)	Literature-grounded
Multiple stable states; hysteresis	Holling (1973); Scheffer (2009)	Hypothesis for this domain
All thresholds, delays and rates	—	Unspecified; would be swept in any build

The share of hypotheses and assumptions is high, and that is deliberate. The document is structured so that audit offices, civil-service reformers and public-administration scholars can move each mechanism from hypothesis towards evidence; it is built to invite that expertise, not to substitute for it.

A.2 What the Document Does and Does Not Claim

It claims that CPAF's seven lenses organise the questions and mechanisms of this domain. It proposes mechanisms, hypothesises configurations, sets out a candidate constraint with two rivals, and identifies readings (direction of knowledge flow, outsourced share of analytical work, command of the method) that spend and headcount omit. It does not claim that any configuration is stable, that the strict gate holds, or that any result has been computed. It does not predict whether a given administration will rebuild its capacity, rank administrative traditions, judge any reform programme, firm or profession, or prescribe a timeline; any such reading is a misuse. It treats consultants as a legitimate and often necessary source of bounded expertise; its subject is structural role, not the competence or integrity of individuals.

A.3 How This Document Was Produced

The domain analysis was developed in a working dialogue with an AI system (Claude), seeded with the CPAF approach and the author's essay, under the author's direction. The author defined the questions, decided what was kept, and is responsible for the result. Three rules governed the work.

Every mechanism was anchored in the published literature before it was included, and its status (empirical, mechanism-grounded or hypothesis) is declared in the table in A.1. Evidence that cuts against the argument is reported alongside evidence that supports it. Every reference was checked against the source or the publisher's record, and any attribution that did not match the record was corrected before inclusion.

The document is itself an example of the compensating use of AI described in Section 8.4: the AI system extended what a single analyst could do, and it remained an instrument rather than a substitute because the author framed the questions and checked every mechanism and reference.

Annex B. References

Entries marked ✓ were checked against the publisher or a repository record, or against the source itself, while preparing this document. The one unmarked entry is the author's own dissertation.

Consultocracy, Policy Capacity and the Hollow State

Andrews, R., & Beynon, M. J. (2024). The revolving door in UK government departments: A configurational analysis. Regulation & Governance, 18(2), 590–611. https://doi.org/10.1111/rego.12544 ✓

Asai, K., Kawai, K., & Nakabayashi, J. (2021). Regulatory capture in public procurement: Evidence from revolving door bureaucrats in Japan. Journal of Economic Behavior & Organization, 186, 328–343. https://doi.org/10.1016/j.jebo.2021.03.031 ✓

Howlett, M., & Migone, A. (2013). Policy advice through the market: The role of external consultants in contemporary policy advisory systems. Policy and Society, 32(3), 241–254. ✓

Kirkpatrick, I., Sturdy, A. J., Reguera Alvarado, N., & Veronesi, G. (2023). Beyond hollowing out: Public sector managers and the use of external management consultants. Public Administration Review, 83(3), 537–551. https://doi.org/10.1111/puar.13612 ✓

Marciano, R. (2024). Beyond consultocracy and servants of power: Explaining the role of consultants in policy formulation. Governance, 37(1), 39–58. https://doi.org/10.1111/gove.12743 ✓

Mazzucato, M., & Collington, R. (2023). The Big Con: How the Consulting Industry Weakens Our Businesses, Infantilizes Our Governments and Warps Our Economies. New York: Penguin Press. ✓

Seabrooke, L., & Sending, O. J. (2022). Consultancies in public administration. Public Administration, 100(3), 457–471. https://doi.org/10.1111/padm.12844 ✓

Sturdy, A., Wright, C., & Wylie, N. (2015). Management as Consultancy: Neo-bureaucracy and the Consultant Manager. Cambridge: Cambridge University Press. ✓

Veselý, A. (2013). Externalization of policy advice: Theory, methodology and evidence. Policy and Society, 32(3), 199–209. ✓

Wu, X., Ramesh, M., & Howlett, M. (2015). Policy capacity: A conceptual framework for understanding policy competences and capabilities. Policy and Society, 34(3–4), 165–171. https://doi.org/10.1016/j.polsoc.2015.09.001 ✓

Ylönen, M., & Kuusela, H. (2019). Consultocracy and its discontents: A critical typology and a call for a research agenda. Governance, 32(2), 241–258. https://doi.org/10.1111/gove.12369 ✓

Capabilities, Learning and Public Value

Bozeman, B. (2007). Public Values and Public Interest: Counterbalancing Economic Individualism. Washington, DC: Georgetown University Press. ✓

Cohen, W. M., & Levinthal, D. A. (1990). Absorptive capacity: A new perspective on learning and innovation. Administrative Science Quarterly, 35(1), 128–152. https://doi.org/10.2307/2393553 ✓

DiMaggio, P. J., & Powell, W. W. (1983). The iron cage revisited: Institutional isomorphism and collective rationality in organizational fields. American Sociological Review, 48(2), 147–160. https://doi.org/10.2307/2095101 ✓

Harvey, G., Jas, P., Walshe, K., & Skelcher, C. (2010). Absorptive capacity: How organisations assimilate and apply knowledge to improve performance. In K. Walshe, G. Harvey & P. Jas (Eds.), Connecting Knowledge and Performance in Public Services (ch. 11, pp. 226–250). Cambridge: Cambridge University Press. ✓

Kattel, R., & Mazzucato, M. (2018). Mission-oriented innovation policy and dynamic capabilities in the public sector. Industrial and Corporate Change, 27(5), 787–801. https://doi.org/10.1093/icc/dty032 ✓

Mazzucato, M., & Kattel, R. (2026). Market-shaping states: A new theory of public sector capacities and capabilities. UCL Institute for Innovation and Public Purpose, Working Paper Series (IIPP WP 2026-03). ✓

Moore, M. H. (1995). Creating Public Value: Strategic Management in Government. Cambridge, MA: Harvard University Press. ✓

Polanyi, M. (2009). The Tacit Dimension (with a new foreword by A. Sen). Chicago: University of Chicago Press. (Original work published 1966) ✓

Evidence, Narrative and Policy Process

Jones, M. D., Smith-Walter, A., McBeth, M. K., & Shanahan, E. A. (2023). The Narrative Policy Framework. In C. M. Weible (Ed.), Theories of the Policy Process (5th ed., pp. 161–195). New York: Routledge. ✓

Parkhurst, J. (2017). The Politics of Evidence: From Evidence-Based Policy to the Good Governance of Evidence. Abingdon: Routledge. ✓

Strassheim, H., & Kettunen, P. (2014). When does evidence-based policy turn into policy-based evidence? Configurations, contexts and mechanisms. Evidence & Policy, 10(2), 259–277. https://doi.org/10.1332/174426514X13990433991320 ✓

Complexity, Leverage and Validation

Abboudi, B. (2026). Modelling Policy Complexity: A Computational Framework for Analysing Multi-Ministry Coordination. Master's dissertation, Jindal School of Government and Public Policy, O.P. Jindal Global University.

Barlas, Y. (1996). Formal aspects of model validity and validation in system dynamics. System Dynamics Review, 12(3), 183–210. https://doi.org/10.1002/(SICI)1099-1727(199623)12:3<183::AID-SDR103>3.0.CO;2-4 ✓

Holling, C. S. (1973). Resilience and stability of ecological systems. Annual Review of Ecology and Systematics, 4, 1–23. ✓

Meadows, D. (1999). Leverage Points: Places to Intervene in a System. Hartland, VT: The Sustainability Institute. ✓

Scheffer, M. (2009). Critical Transitions in Nature and Society. Princeton, NJ: Princeton University Press. ✓

Reports, Guidance and Data

Consultancy UK. (2025, March 7). Civil servants expect consulting spending to rise. ✓

Consultancy UK. (2026, May 26). Government spending on consulting grows to record £3.7 billion. ✓

Digital, Data and Technology Profession (2023). Intellectual Property Rights Guidance Note (The Digital, Data and Technology Playbook). London: GOV.UK, 20 June 2023. ✓

National Audit Office (2025). Lessons Learned: The Government's Use of External Consultants. HC 1381, Session 2024–2026, 21 November 2025. London: National Audit Office. ✓

Senate Finance and Public Administration References Committee (2024). Management and Assurance of Integrity by Consulting Services: Final Report, June 2024. Canberra: Commonwealth of Australia. ✓

End of working document.

Prepared by Bilel ABBOUDI as a conceptual proof of concept. A working paper, not a forecast, an assessment of any administration, or procurement advice; open to discussion.

CPAF v6 — Computational Policy Analysis Framework.